

VIRTUAL STAR NETWORK*ensa*
Background and Summary of the Invention

The present invention relates to wavelength division multiplexing (WDM) optical communication networks, and particularly to a WDM optical ring network that implements a virtual star network. More particularly, the present invention relates to a WDM optical network including a plurality of terminal nodes that each send and receive signals at a particular wavelength and a head-end node that switches the signals between the terminal nodes to implement a virtual star network over an optical ring.

A transition from shared media networks to switched data networks has revolutionized the Local Area Network (LAN) arena in recent years. This transition has taken place for virtually all LAN technologies, including Ethernet, Token Ring and FDDI (Fiber Distributed Data Interface). A protocol that has achieved some success in data network applications is ATM (Asynchronous Transfer Mode), which is inherently a switched technology. This change to switched networks has occurred quickly because of advantages over shared media networks, such as an increase in the bandwidth available to each terminal without increasing the complexity of the terminal.

Each LAN terminal now typically has dedicated bandwidth, which is not affected by the presence of other terminals in the network. Each terminal can be assigned a rate suitable for its needs (e.g., 10 megabits per second (Mbps) for a desktop device and 100 Mbps for a server). Different rate terminals can be attached to the same network, and changing the rate for one terminal does not affect other terminals on the network.

Unlike shared media networks, switched networks imply the existence of a centralized switching element to which all terminals must be connected. One reason the quick migration to switched environments was possible was because the wiring topology in most corporate networks had already been converted to a star topology in order to take advantage of the previous generation of networking hardware, that is, shared hubs.

Telecommunication carriers would benefit from increased capacity and upgrade flexibility associated with switched environments. Current generation local

exchange transmission networks are usually based on fiber rings using the Synchronous Optical Network (SONET) standard or the Synchronous Digital Hierarchy (SDH) standard (SONET is the North American version of the international SDH standard).

Most telecommunication carriers have a two-level hierarchy of rings.

- 5 The first level rings, sometimes referred to as Metro rings or Inter-Office Facilities (IOFs), usually use OC-48 rates (2.488 gigabits per second (Gbps)) for very high speed, wide area networking. The second level rings, sometimes referred to as Access rings, link access devices (e.g., a Digital Subscriber Line Access Multiplexer) for eventual connection to a Metro ring and use OC-3 (155.52 Mbps) or OC-12 (622.08
- 10 Mbps) rates. As capacity requirements increase, both ring types will need to migrate to higher speeds.

- Applying LAN terminology, SONET rings should be classified as "shared media" environments. Devices are connected to the SONET rings using Add/Drop Multiplexers (ADMs), and the total bandwidth of the ring (OC-3/12/48) is
- 15 shared among all nodes on the ring (in this analogy, ADMs on the SONET rings are regarded as terminals in the network). When the required bandwidth exceeds the ring capacity, all ADMs on the ring need to be upgraded to a higher speed.

- The SONET/SDH equivalent of a switched network would be to use a terminal multiplexer at each node and to connect all of them to a SONET cross-
- 20 connect system. This, however, would require the physical fiber topology to be star-like, similar to the wiring structures of corporate LANs. There are, however, two significant differences between the LAN environment and telecommunication transmission networks. First, the distances covered by the telecommunication network are much larger. Second, reliability is of much greater importance (loss of connectivity
- 25 to a LAN terminal usually affects a single user, but loss of connectivity to an ADM on a SONET ring may affect thousands).

- These two aspects of local exchange transmission networks make rings a preferred physical topology for telecommunication applications. A physical ring topology achieves lower total fiber mileage compared to connecting the nodes in a
- 30 mesh or star topology. Moreover, the existence of two alternate paths between each pair of nodes is an inherent property of a ring topology. The present invention implements a logical star network over a physical ring topology to retain the physical

benefits of the ring topology while providing switched data network advantages of the star topology.

According to an aspect of the present invention, a system for communicating between a plurality of nodes coupled to an optical wavelength division multiplexed ring network is provided. The system includes at least two terminal nodes and a head-end node. Each terminal node has a communication subsystem, a tributary subsystem, and a multiplexing subsystem. A first terminal node's communication subsystem is configured to be coupled to the ring network to receive and to transmit signals at a first wavelength and to permit signals at other wavelengths to pass. A second terminal node's communication subsystem is configured to be coupled to the ring network to receive and to transmit signals at a second wavelength and to permit signals at other wavelengths to pass. The tributary subsystems are configured to be coupled to a plurality of devices to enable the devices to communicate over the ring network. The multiplexing subsystems are coupled to the tributary subsystem and to the communication subsystem to channel signals between the plurality of devices and the ring network. The head-end node is coupled to the ring network to receive and to transmit signals at both the first and second wavelengths. The head-end node has a demultiplexer to isolate signals received at the first and second wavelengths, an integral cross-connect module to determine an output wavelength at which to transmit received signals, and a multiplexer to combine the received signals for transmission on the ring network at the first and second wavelengths.

According to an aspect of preferred embodiments, the terminal node communication subsystems include an optical add/drop multiplexer coupled to the ring network. The head-end node can include a tributary subsystem configured to be coupled to a plurality of devices to enable the devices to communicate over the ring network. The terminal nodes and head-end node can be configured to receive and to transmit signals using a synchronous optical network communication standard such as SONET/SDH.

According to other aspects of preferred embodiments, the head-end node receives and transmits signals using a synchronous optical network communication standard, such as SONET/SDH. A subset of the signals further use a communication protocol, such as ATM or IP, framed by the communication standard.

The head-end node includes at least one protocol subsystem to determine address information for the communication protocol, and the head-end node is configured to send signals using the communication protocol to the at least one protocol subsystem.

According to yet other aspects preferred embodiments, the head-end node receives and transmits signals using a synchronous optical network communication standard, such as SONET/SDH. A subset of the signals further use a communication protocol, such as ATM, framed by the communication standard and a second subset of the signals further use a second communication protocol, such as IP, framed by the communication standard. The head-end node includes first and second protocol subsystems to determine address information for the first and second communication protocols. The head-end node is configured to send signals using the first communication protocol to the first protocol subsystem and to send signals using the second communication protocol to the second protocol subsystem.

According to still other aspects of preferred embodiments, the head-end node includes first and second transmitters coupled to the multiplexer to send signals at the first and second wavelengths, respectively. The head-end node includes first and second receivers coupled to the demultiplexer to receive signals at the first and second wavelengths, respectively.

According to yet still other aspects of preferred embodiments, the ring network includes a first ring for transmitting information in a clockwise direction and a second ring for transmitting information in a counter-clockwise direction. The communication subsystems include a pair of transceivers coupled to the first and second rings, respectively. The demultiplexers include a pair of demultiplexers coupled to the first and second rings, respectively. The multiplexer includes a pair of multiplexers coupled to the first and second rings, respectively. The communication subsystem can further include a selector that compares a pair of signals received by the pair of transceivers and selects a signal from the pair of signals based on a quality parameter of each signal. The head-end node can further include a selector that compares a pair of signals received by the pair of demultiplexers and selects a signal from the pair of signals based on a quality parameter of each signal.

According to another aspect of the present invention, a system for communicating between a plurality of nodes coupled to an optical wavelength division

[illegible]

According to aspects of preferred embodiments, two terminal nodes with communication subsystems as discussed above can be provided. Each communication subsystem further includes a pair of transceivers coupled to the first and second rings, respectively. Each communication subsystem includes a selector to select a signal from the pair of signals received by the pair of transceivers based on a quality parameter of each signal.

According to still another aspect of the invention, a system for communicating between a plurality of nodes coupled to an optical wavelength division multiplexed ring network using a primary communication standard, such as SONET/SDH, is provided. At least some of the nodes in the network send and receive signals using at least one secondary communication protocol, such as ATM or IP, framed by the primary communication standard. The system includes a head-end node coupled to the ring network to receive and to transmit signals at first and second wavelengths. The head-end node has a demultiplexer, cross-connect, multiplexer, and at least one protocol subsystem. The demultiplexer isolates signals received at the first and second wavelengths. The cross-connect module to determine an output wavelength at which to transmit received signals. The multiplexer combines the isolated received signals to transmit on the ring network at the first and second wavelengths. The protocol subsystem is coupled to the cross-connect module to determine address information for the at least one secondary protocol of the received

signals. The secondary communication protocol can be a protocol encapsulated within a protocol, such as IP encapsulated within ATM.

According to yet another aspect of the invention, a method of communicating signals in an optical wavelength division multiplexed ring network between two nodes using a primary communication standard and a secondary communication protocol framed by the primary communication standard is provided. The method includes providing a first terminal node coupled to the ring network to receive and to transmit signals at a first wavelength and to permit signals at other wavelengths to pass. A second terminal node is provided for coupling to the ring network to receive and to transmit signals at a second wavelength and to permit signals at other wavelengths to pass. A head-end node is provided for coupling to the ring network to receive and to transmit signals at both the first and second wavelengths. Signals transmitted at the first and second wavelengths are received at the head-end node. Destination address information for signals received at the head-end node is determined based on information in the signal encoded by the secondary communication protocol. Signals received at the head-end node at one of the first and second wavelengths are retransmitted based on the destination address information.

Additional features of the invention will become apparent to those skilled in the art upon consideration of the following detailed description of the preferred embodiment exemplifying the best mode of carrying out the invention as presently perceived.

Brief Description of the Drawings

Fig. 1 is a logical network diagram showing a head-end node coupled to a WDM optical ring network for communication with five terminal nodes at five different wavelengths to implement a virtual star network;

Fig. 2 is a block diagram showing the architecture of a terminal node including an optical subsystem coupled to a ring network having a pair of counter-rotating fiber rings, a multiplexing subsystem coupled to the optical subsystem, a tributary subsystem coupled to the multiplexing subsystem for connecting client systems to the ring network, and a control subsystem coupled to the optical, multiplexing, and tributary subsystems;

Fig. 3 is a block diagram showing an architecture for an alternative terminal node according to the present invention using the same optical and control subsystems of the terminal node of Fig. 2, but in which the functionality of the multiplexing and tributary sub-systems is assumed to be carried out by an external system;

Fig. 4 is a block diagram showing an embodiment of the optical subsystem of Figs. 2 and 3 including a pair of optical add/drop multiplexers coupled to the fiber rings and a pair of WDM transceivers coupled to the multiplexing subsystem;

Fig. 5 is a conceptual diagram showing the functionality of the optical add/drop multiplexers of Fig. 4;

Fig. 6 is a block diagram showing an architecture of a head-end node somewhat similar to the architecture of the terminal node of Fig. 2, including an optical subsystem coupled to the ring network, a protocol subsystem coupled to the optical subsystem, a tributary subsystem coupled to the protocol subsystem, and a control subsystem coupled to the optical, protocol, and tributary subsystems;

Fig. 7 is a block diagram showing an embodiment of the optical subsystem of Fig. 6 including a pair of WDM demultiplexers coupled to the fiber rings to receive signals, a pair of WDM multiplexers coupled to the fiber rings to transmit signals, and pairs of transceivers dedicated to each terminal wavelength in the optical WDM network coupled to the multiplexers, demultiplexers, and to the protocol subsystem;

Sub B. 7 Figs. 8a and 8b are conceptual diagrams showing the functionality of the WDM demultiplexers and multiplexers, respectively, of Fig. 7;

Fig. 9 is a conceptual block diagram showing network reliability aspects of a virtual star network according to the present invention;

Fig. 10 is block diagram showing signal flows through a terminal node according to the present invention;

Fig. 11 is a block diagram showing signal flows through a head-end node according to the present invention having an internal cross-connect that can be a SONET/SDH cross-connect matrix, an ATM switch, a VPX matrix, or an IP routing matrix;

Fig. 12 is a block diagram showing signal flows through an alternative head-end node embodiment according to the present invention configured to route both ATM and IP protocols over a SONET/SDH-based virtual star network; and

Fig. 13 is a block diagram somewhat similar to Fig. 12 showing signal flows through another alternative head-end node configured to route IP data encapsulated in ATM that in turn is transmitted over a SONET/SDH-based network.

Detailed Description of Drawings

A virtual star network 20 according to the present invention implements a virtual star over a physical fiber ring 22 by using a head-end node 24 that provides switching functionality for a plurality of terminal nodes 26 as shown in Fig. 1. Fiber ring 22 illustratively is a pair of counter-rotating fibers 23, 25 as shown, for example, in Figs. 2-4, to provide a redundant, bidirectional communication path so that virtual star network 20 provides a redundant, bidirectional network architecture.

As discussed in more detail below, head-end and terminal nodes 24, 26 provide for coupling client systems (not shown) to ring 22 for communication over virtual star network 20. Each terminal node 26 has a separate communication channel 28 over network 20 to head-end 24 but does not have a direct communication channel to any other terminal node 26. Each terminal node 26 sends all network traffic from its connected client systems to head-end 24, which effects a cross-connection function and sends to each terminal node 26 the network traffic intended for it. Network 20 thus implements a virtual star over a physical ring, with head-end node 24 as the star's logical center, or hub, and terminal nodes 26 as the star's logical points.

Implementing a virtual star over fiber ring 22 requires the ability to transmit multiple and separate logical channels 28 within fiber ring 22. An illustrative technology that provides this capability is Dense Wavelength Division Multiplexing (DWDM), which allows the use of multiple independent optical channels 28, each at a different wavelength, within the same optical fiber. All channels 28 are illustratively located in the 1550 nm band. Each terminal node 26 is assigned a specific wavelength to simplify the optical communication hardware, although designs that allow for changing or switching between wavelengths are known and can be used. Each terminal node 26 uses its specific wavelength as an "uplink channel" to transmit its

traffic to head-end 24. Head-end 24 uses the same wavelength as a "downlink channel" to transmit the return traffic intended for each terminal node 26.

Transition of a SONET/SDH ring 22 to a virtual star, switched architecture improves the cost-effectiveness and network capacity of an optical ring network. For example, with a conventional OC-48 ring, each node on the ring requires an OC-48 ADM. Each such ADM must be able to handle the full bandwidth of the ring and must be able to drop and to add tributaries into this high-speed channel. The aggregate bandwidth available to all nodes on the ring is about 2.5 Gbps, which must be shared among all nodes. It follows that the aggregate rate available to client systems coupled to each node is much less than the ring bandwidth; it is usually limited to about 622 Mbps (OC-12).

An approach based on a switched, virtual star architecture according to the present invention, is as follows. The OC-48 ADM at each node is replaced by an OC-12 terminal multiplexer, provided with an optical sub-system as discussed below that provides the functionality of a terminal node 26 in the virtual star network. A head-end node 24 is provided to terminate all uplinks and downlinks to and from terminal nodes 26 to effect a SONET/SDH cross-connect function for the traffic from the terminals.

Virtual star network 20 provides the functional behavior of an ADM-based OC-48 ring, but with added benefits. For example, because each terminal node 26 has available to it the full bandwidth for its wavelength, the aggregate rate available to each node in the virtual star is 622 Mbps (OC-12), regardless of the number of nodes in the network. For comparison, in an ADM-based OC-48 ring network of more than 8 nodes, the average capacity available to each node typically is less than 622 Mbps. Network transceiver hardware for terminal nodes 26 in virtual star network 20 is significantly less complex than that of an OC-48 ADM. This is due to the lower rates handled by terminal nodes 26 and the simpler functionality of a SONET terminal multiplexer as compared to an ADM. Less complexity results in lower cost and higher software and hardware reliability.

In virtual star network 20, when the traffic demands of a specific terminal node 26 exceed an aggregate of OC-12 terminal multiplexer, then that node can be upgraded (e.g., to an OC-48 Terminal Mux). This upgrade, however, need not

affect the other nodes in the ring and will not as a practical matter decrease the bandwidth available to them. Thus, virtual star network 20 provides for lower cost in addition to improved functional benefits such higher reliability and improved effective throughput.

5 A virtual star architecture can provide higher effective throughput at lower cost for SONET/SDH rings at other rates. For example, a virtual star network 20 with OC-3/STM-1 terminal nodes 26 and a suitable head-end node 24 can replace an OC-12/STM-4 ADM-based SONET/SDH ring. Similarly, a virtual star network 20 with OC-48/STM-16 terminal nodes 26 and a suitable head-end 24 can replace an
10 OC-192/STM-16 ADM-based SONET/SDH ring.

 The virtual star architecture is applicable to rings other than SONET/SDH rings. It can be used, for example, to enhance and simplify an ATM-based transmission network (ATM Ring). Each OC-48 ATM ADM can be replaced by a simpler OC-12 ATM multiplexer. The head-end node 26 in this case would need
15 to incorporate ATM VPX functionality. The benefits discussed above for SONET/SDH networks apply equally to the ATM Ring. In addition, virtual star ATM networks will achieve lower transmission delay, and in particular lower call-setup delay, when compared to a traditional ATM Ring, because the virtual star architecture guarantees two hops between terminal nodes as compared to N hops for a traditional
20 ring.

 A virtual star network 20 according to the present invention can be used for the hubbing of IP (Internet Protocol) traffic in a ring network 22 to a centralized backbone router. Multiple IP sources can be concentrated by a terminal node 26 to a single high-speed IP link or channel, and transported to the head-end 24
25 over the associated wavelength. In this application, head-end node 24 incorporates the functionality of an IP router or IP switch and processes the traffic from the multiple sources.

 Because the virtual star architecture allows the rate and the protocol used by each terminal node 26 to be independent of those associated with other
30 terminal nodes 26, it allows the deployment of mixed-protocol rings, whereby, for example, some terminals run SONET/SDH and others run ATM, and still others IP. The head-end node 24 can provide for SONET/SDH cross-connect on all

SONET/SDH channels, an ATM VPX on all ATM channels, as well as IP routing for all IP channels. In an alternative embodiment discussed below, IP can be embedded within ATM and then transmitted over a SONET/SDH-based network.

Each terminal node 26 includes an optical subsystem 30, a multiplexing subsystem 32, a tributary subsystem 34, and a control subsystem 36 as shown in Fig. 2. Optical subsystem 30 implements an optical add/drop function that provides for transmission of an aggregated signal produced by multiplexing subsystem 32 on its dedicated optical channel 28 to head-end node 24. Tributary subsystem 34 includes interface cards of different types and rates to provide client interfaces 38 as appropriate for a given application, to which client systems (not shown) are coupled as desired. A communication interface 40 between tributary subsystem 34 and multiplexing subsystem 32 is provided to insulate multiplexing subsystem 32 from the effects of interfacing different client systems to tributary subsystem 34.

Multiplexing subsystem 32 receives the different tributary channels or signals from client systems coupled to tributary subsystem 34 over interface 40 and aggregates them onto a single stream with a format appropriate to the desired application, such as SONET/SDH, ATM or IP. This single stream is then passed over interface 42 to optical subsystem 30, which in turn transmits it over ring 22 to head-end node 24. Conversely, multiplexing subsystem 32 receives an incoming aggregated signal stream from optical subsystem 30 that multiplexing subsystem 32 then demultiplexes into its constituent tributaries, each of which is then routed over interface 40 to a corresponding interface card in tributary subsystem 34.

^{Sub 32} An alternative terminal node 26' architecture is provided without a tributary subsystem 34 as shown in Fig. 3. An external system 42 such as a SONET/SDH terminal mux for a SONET/SDH application, an ATM Service Mux for an ATM application, or an IP router for an IP application, is used to replace tributary subsystem 34 and part of multiplexing subsystem 32 in terminal node 26. A modified multiplexing subsystem 32' with an interface 40' to external subsystem 34 is provided, with the optical subsystem 30 and interface 42 being the same as the embodiment of Fig. 2.

Control subsystem 36 manages, configures and monitors the operation of the optical, multiplexing, and tributary subsystems 30, 32, 34. The control and

optical subsystems 36, 30 are common to all applications of terminal node 26 in virtual star network 20, whereas the tributary and multiplexing subsystems 32, 34 are application specific. Tributary and multiplexing subsystems 32, 34 are implemented by plug-in modules that are accommodated by a common platform that includes control and optical subsystems 36, 30.

An illustrative architecture for optical subsystem 30 of terminal node 26 includes a pair of DWDM transceiver modules 46 coupled to a respective pair of optical add/drop multiplexers (OADMs) 48 as shown in Fig. 4. Fibers 23, 25 carry an aggregate DWDM signal 50 illustratively containing four constituent wavelengths 51, 52, 53, 54 as shown in Fig. 5. OADM 48 drops a specific wavelength from a DWDM signal on fiber 23, 25 and routes it to DWDM transceiver module 46. Similarly, an optical signal generated by DWDM transceiver module 46 (and having the same wavelength) is inserted by OADM 48 into the aggregate DWDM signal 50. Each OADM 48 for terminal nodes 26 is specific to a given wavelength and will pass all other wavelengths through transparently. OADM components are commercially available from several vendors, such as a model ADOM200031310 available from E-TEK.

DWDM transceiver module 46 includes a transmitter and a receiver (not shown) for a specific wavelength λ . The transmitter transforms an electrical signal (generated by multiplexing subsystem 32) to an optical signal at wavelength λ . The receiver transforms an optical signal to an electrical signal provided to multiplexing subsystem 32. Suitable transmitters and receivers for DWDM transceiver module 46 are commercially available from several vendors, such as models HFCT-10XX, from Hewlett-Packard (transmitter), and RGR-2622, also from Hewlett-Packard.

For protection purposes, two copies of the optical signal are received at terminal node 26, one from each fiber 23, 25. The dual set of OADMs 48 and DWDM transceiver modules 46 allows terminal node 26 to monitor both incoming signals for quality and select the better one at each point in time for protocol processing. This monitoring and selection is carried out by the multiplexing subsystem 32, although this function can alternatively be provided by optical subsystem 30 or by a separate subsystem. Terminal node 26 similarly transmits two identical copies of its outgoing

signal, one for each fiber 23, 25, to provide for redundancy on ring 22 and at head-end node 24. This redundant protection scheme is discussed in more detail below.

Head-end node 24 includes an optical subsystem 56, protocol subsystem 58, tributary subsystem 60, and control subsystem 62 as shown in Fig. 6.

- 5 Optical subsystem 56 terminates and generates all the terminal uplinks and downlinks, respectively, for head-end node 24. As discussed below, optical subsystem 56 incorporates multiplexing and demultiplexing functionality for the DWDM channels, as well as suitable receivers and transmitters.

- 10 Protocol subsystem 58 provides protocol-related functions such as the cross-connect/switching function and all protocol-specific processing. When head-end node 24 is used in a SONET/SDH application, protocol subsystem 58 provides the functionality of a SONET/SDH cross-connect, as well as all SONET/SDH-related protocol processing. When head-end node 24 is used in an ATM application, protocol subsystem 58 provides the functionality of an ATM switch or VPX, along with the
- 15 associated protocol processing. When head-end node 24 is used for an IP application, protocol subsystem 58 provides the functionality of an IP router or IP switch, along with the associated protocol processing.

- Similar to tributary subsystem 34 of terminal nodes 26, head-end tributary subsystem 60 allows head-end node 24 to have local client systems (not shown) and is coupled to protocol subsystem 58 by an interface 64. Head-end node 24 can accommodate various client system hardware without affecting protocol subsystem 58. Head-end node 24 can perform a hubbing function for the terminal nodes 26, in which all traffic generated by terminal nodes 26 is dropped as local tributaries of head-end node 24. The tributary capacity of head-end node 24 could be equal to the
- 25 aggregate capacity of the traffic of all terminals 26.

- Again similar to the architecture of terminal nodes 26, head-end control subsystem 62 manages, configures and monitors the operation of head-end optical, protocol, and tributary subsystems 56, 58, 60. Control and optical subsystems 62, 56 are common to all applications of head-end node 26 in virtual star network 20,
- 30 whereas tributary and protocol subsystems 60, 58 are application specific. Tributary and multiplexing subsystems 60, 58 are implemented by plug-in modules that are accommodated by a common platform that includes control and optical subsystems 62,

56. Head-end node 24 provides a single, modular package that includes both the optical and electronic components for implementing virtual star network 20 in a single, integral package.

Head-end optical subsystem 56 includes a pair of DWDM mux/demux modules 68 as shown in Fig. 7. DWDM mux/demux modules 68 have a multiplexer 70 and a demultiplexer 72. Multiplexer 70 multiplexes several optical signals 71, each of a different wavelength, into a single DWDM fiber signal output 71' as shown in Fig. 8a. Demultiplexer 72 separates a DWDM signal 71' carried on a fiber into separate wavelength outputs 73 as shown in Fig. 8b. DWDM multiplexers and demultiplexers are commercially available from several vendors, such as model DWDM2F0821310 (multiplexer) and DWDM2F0822310 (demultiplexer), both from E-TEK.

Head-end optical subsystem 56 further includes pairs of DWDM transceiver modules 74 that each have a transmitter/receiver pair 76 for each specific wavelength λ_i in DWDM signal 50. The transmitter transforms an electrical signal provided from protocol subsystem 58 over interface 66 to an optical signal at wavelength λ_i provided to DWDM mux/demux module 68. The receiver transforms an optical signal at wavelength λ_i from mux/demux module 68 to an electrical signal provided to protocol subsystem 58. Such transmitters and receivers are commercially available from several vendors, such as models HFCT-10XX (transmitter) and RGR-2622 (receiver), both from Hewlett-Packard.

For protection purposes, two copies of each optical signal are received at head-end node 24, one on each fiber 23, 25. The dual set of transceivers per wavelength allows head-end node 24 to monitor each pair of corresponding incoming signals for quality and to select the better one at each point in time for protocol processing. Similarly, head-end node 24 transmits two identical copies of each signal, one of each fiber 23, 25, for the benefit of the corresponding terminal node 26. This monitoring and selection is carried out by protocol subsystem 58 as discussed in more detail below.

The above-described implementation head-end and terminal nodes 24, 26 for virtual star network 20 provide for high network reliability. For example, virtual star network 20 is protected from failures due to a "fiber-cut", where the physical fiber pair 23, 25 becomes disconnected at some point. This could be the

result of a fiber physically being cut or of a malfunctioning of some equipment, and would not affect operation of remaining fiber or equipment. Similarly, virtual star network 20 will operate notwithstanding a transmission impairment that renders a segment of ring 22 unusable for one or more wavelengths. Furthermore, the failure or degradation in the quality of a transmitter/receiver in either head-end 24 node or a terminal node 26 will not impair performance of virtual star network 20.

Virtual star network 20 provides a "1+1" protection approach, similar in spirit to that used in SONET UPSR (Unidirectional Path Switched Ring) rings. Each terminal node 26 transmits its signal in both directions on ring 22. Likewise, head-end node 24 transmits all its signals in both directions. This is achieved by using the two transmitters available in each terminal node 26 and the two transmitters per wavelength in head-end node 24. As discussed above, both terminal nodes 26 and head-end node 24 also have two receivers (two per wavelength in the head-end). Thus the two incoming signals can be monitored simultaneously for quality, and the best-received signal can be selected at every point in time.

Virtual star network 20 provides for fail-safe operation as shown in Fig. 9. For a virtual star network 20 with two terminal nodes 26 and six fiber segments 78, 79, 80, 81, 82, 83, if a fiber cut or any transmission impairment renders any segment unusable, each terminal node 26 can still reach head-end node 24, and vice versa. Similarly, a failure of either a transmitter or a receiver in either head-end node 24 or a terminal node 26 will effect the same result as a fiber cut. Performance of virtual star network 20 will be unaffected because of the redundant copy of the signal.

Referring to Fig. 10, terminal node 26 provides for interfacing client systems to virtual star network 20 for either protocol-specific or mixed protocol applications. Signals from all client systems coming into tributary subsystem 34 are multiplexed into a single high-speed signal 84 in multiplexing subsystem 32. The multiplexed signal 84 is duplicated and passed on to both DWDM transceivers 46.

In receiving data over fibers 23, 25, multiplexing subsystem 32 selects the better of the two signals 86 received from transceivers 46 and breaks the selected signal up into its constituent tributary signals. This selection is made in real-time so that the best available signal is always used, and is based on signal quality parameters available from the transmission standard and network protocol being used. For

example, as discussed below, overhead parity bytes are used for the case of the SONET/SDH standards. Any signal quality parameter can be used, information encoded on signals received from fibers 23, 25 as well as signal quality parameters generated by terminal node 26, such as via monitoring and comparison of signals. The
5 selected signal is demultiplexed and sent over interface 40 to tributary subsystem 34.

Tributary subsystem 34 contains tributary cards 88 that support client systems of various rates, such as, for example, T1, T3, OC-3, etc. (for SDH systems, E1, E3, STM-1, etc.). The number of tributary cards 88, as well as the number of client systems supported by each such card, is arbitrary, so long as the aggregate
10 capacity of all the tributary cards 88 does not exceed the rate of the high-speed channel 28.

In a SONET/SDH application, the functionality of multiplexing and tributary subsystems 32, 34 in terminal node 26 is equivalent to a SONET/SDH terminal multiplexer that supports 1+1 line redundancy. High speed multiplexed signal
15 84 that is passed to DWDM transceivers 46 can be, for example OC-12/STM-4 or OC-48/STM-16. The selection of the better of the two signals 86 received from transceivers 46 is based upon signal quality parameters for signals 86 as defined by the SONET/SDH standards, such as the overhead parity bytes, bit error rates, etc.

All non-SONET/SDH tributary signals from client systems (e.g., T1, T3
20 or E1, E3) undergo mapping in tributary subsystem 34 to SONET/SDH, according to the SONET/SDH standards. All the resulting SONET/SDH tributary channels are then multiplexed into the high-speed SONET/SDH channel 84 in multiplexing subsystem 32. Various SONET/SDH terminal multiplexers that provide this functionality are commercially available, such as model FLM 600 from Fujitsu.

25 In an ATM application, the functionality of multiplexing and tributary subsystems 32, 34 generally resembles that of an ATM service multiplexer (ASM) that supports 1+1 line redundancy. In other words, the traffic from all incoming client systems is converted to ATM cells if it was not ATM to begin with. The ATM cells are then multiplexed into a single ATM high-speed signal 84 (the different tributary
30 signals are associated with different VPCs (Virtual Path Connections) or VCCs (Virtual Circuit Connections) in the unified ATM stream). This signal 84 is duplicated and passed on to both WDM transceivers 46 for transmission on fibers 23, 25. In

receiving signals from ring 22, the multiplexing subsystem 32 selects the better of the two incoming high-speed signals and demultiplexes that signal into separate streams intended for client systems coupled to tributary subsystem 34.

As discussed above, tributary cards 88 may support various client system rates, such as T1, T3, OC-3, as well as non-ATM signals such as various LAN technologies (e.g., Ethernet) or audio/video signals. Again, the number tributary cards 88, as well as the number of client systems supported by each card 88, is relatively arbitrary. All non-ATM tributary signals undergo mapping to ATM, according to the ATM standards. All the resulting ATM tributary channels are then multiplexed into the high-speed ATM channel 84.

The receiver function for terminal node 26 used in an ATM application monitors the performance of two incoming signals from fibers 23, 25 by using the SONET/SDH overhead parity bytes, as defined by the SONET/SDH standards (ATM is defined to use SONET/SDH framing and overheads), or by using the ATM OAM (operations and maintenance) performance monitoring provisions, or by using both. The better-received signal 86 is always selected in real time for each incoming signal. The demultiplexer function uses the ATM header on each cell to route each cell to its intended destination tributary card 88, which in turn forward cells to client systems.

Commercially available ATM switches, such as model Lightstream 1010 from Cisco, or or service muxes, such as model 36140 from Siemens, provide the above-described ATM functionality. Terminal node 26 can also implement ATM functionality through commercially available chipsets such as Atlanta from Lucent.

In a pure IP application, the functionality of multiplexing and tributary subsystems 32, 34 resembles that of an IP router, enhanced to support 1+1 line redundancy. In other words, traffic from client systems will be IP traffic, and is multiplexed into a single IP high-speed signal 84 (that is mapped to SONET/SDH using Packet-over-SONET (POS) specifications). Signal 84 is duplicated and passed on to both WDM transmitters. In the receiver direction, the multiplexing subsystem 32 selects the better of the two incoming high-speed signals 86 and demultiplexes the selected signal 86 into separate streams intended for the specific tributary cards 88 and client systems. As discussed above for SONET and ATM applications, tributary cards 88 can support various rates for client systems and support signals for different LAN

technologies (e.g., Ethernet) or other signal types, such as, for example, audio/video signals. Again, the number of tributary cards 88 and supported client systems by each card 88 is relatively arbitrary.

The receiver function in a terminal node 26 in an IP application
5 illustratively monitors the performance of two incoming signals 86 by using the overhead parity bytes in the SONET/SDH frame in which the IP channel is embedded. The better-received signal is used at every point in time. The demultiplexer function uses the IP header on each packet (specifically, the IP address) to route each packet to its intended destination tributary port. Commercially available routers such as the
10 7500 line from Cisco provide this IP functionality. Similarly, commercially available chipsets and software such as AnyFlow 5500 from MMC Networks can be used to achieve the described IP functionality.

Referring now to Fig. 11, tributary subsystem 60 of head-end 24 also contains tributary cards 88 for coupling to client systems of various rates as discussed
15 above for terminal nodes 26. Protocol subsystem 58 includes a selector 90 and a switch matrix 92. Selector 90 provides for selecting signals from fiber rings 23, 25 based on one or more quality parameters as also discussed above. Although illustrated as a component of protocol subsystem 58, selector 90 can be provided as part of optical subsystem 56 or as a logically separate subsystem. Switch matrix 92 provides
20 protocol-specific switching or routing of selected signals between optical subsystem 56 and tributary subsystem 60. As discussed in more detail below, switch matrix 92 can be modified or enhanced to provide for switching or routing multiple protocols.

In a SONET/SDH application, protocol and tributary subsystems 58,
60 of head-end node 24 provide functionality similar to a SONET/SDH cross-connect
25 switch. Switch matrix 92 illustratively provides a SONET/SDH cross-connect matrix that has N input ports and N corresponding output ports. A SONET/SDH signal is present on each input port at one of the rates specified by the SONET/SDH hierarchy (e.g., OC-12, OC-48, etc.). A channel of the same rate is output by matrix 92 on the corresponding output port. Matrix 92 is able to rearrange and remap the sub-channels
30 in the incoming channels to sub-channel slots in the outgoing channels. The remapping function can be configured dynamically. Thus, it is possible to transmit a sub-channel

from any one port to any other port. A multicast capability is also provided by mapping one port to multiple other ports.

Two levels of cross-connecting in matrix 92 are illustratively provided. The first level is STS-1 level cross-connect, in which the granularity of the sub-channels manipulated by matrix 92 corresponds to SONET STS-1 frames (High Order Tributary cross-connecting in SDH terminology). The second level is VT (Virtual Tributary) level cross-connect, in which the granularity of the sub-channels manipulated by the matrix corresponds to SONET Virtual Tributaries (Low Order Tributaries in SDH terminology).

10 In Fig. 11, the input and output ports are associated with either tributary cards 88 or WDM transceivers 76 handling channels originating from terminal nodes coupled to ring 22. Matrix 92 allows a sub-channel originating in either a terminal 26 or in one of the head-end tributary cards 88 to be transmitted to any other (set of) terminal node 26 or tributary card 88. Matrix 92 is configured with sufficient
15 capacity to handle traffic from terminal nodes 26 and client systems coupled to tributary cards 88 without any throughput limitations (blocking).

Designs for SONET/SDH cross-connect matrices are known in the art. Matrix 92 allows for use of virtually any cross-connect design to implement the functionality for protocol subsystem 58 as described above. Similarly, cross-connect
20 systems and electronic components which allow the design of a SONET/SDH cross-connect module are commercially available from various vendors such as PMC-Sierra.

Selector 90 is coupled to dual sets of WDM transceivers 76 (two transceivers per wavelength) that are part of optical subsystem 56 as shown in Fig. 11.
25 Selector 90 illustratively includes discrete selector modules (not shown) for each channel or wavelength. Selector modules plug into a logical block and select the better of two signals per wavelength by monitoring the quality of the two received signals using the provisions in the SONET/SDH header, as explained above. The better received signal is then passed on to the SONET/SDH matrix 92. In the other
30 direction, each signal originated by matrix 92 is duplicated by the appropriate selector module in selector 90 to both transmitters associated with the channel or wavelength.

In an ATM application, the functionality of protocol and tributary subsystems 58, 60 resembles that of an ATM switch or an ATM VP cross-connect (VPX). In this case, matrix 92 is an ATM Switch/VPX module that has N input port and N corresponding output ports. An ATM signal is present on each port at one of the rates supported by matrix 92. A channel of the same rate is output by matrix 92 on the corresponding output port. Matrix 92 routes cells from the incoming channels to output channels based on the VPI/VCI (virtual packet information/virtual circuit information) in its header. In case of a VPX matrix only the VPI field is considered. In case of a full ATM switch matrix both the VPI and the VCI are considered. Matrix 92 contains the required buffering and switching resources.

Matrix 92 input/output ports are associated with either tributary cards 88 or WDM transceivers 76 handling channels originating from terminal nodes 26 on virtual star network 20. Matrix 92 allows a VCC or VPC originating in either any terminal node 26 or local client systems coupled to head-end tributary cards 88 to be transmitted to any other terminal node 26 or local client system.

Again, designs of ATM switches and VPX matrices are known in the art and such known designs can be used for matrix 92 in an ATM application. ATM switch/VPX systems and electronic components that allow for implementation of a suitable ATM switch/VPX matrix 92 are also commercially available from various vendors such as Lucent.

Selector 90 in ATM applications monitors the quality of the two received signals per wavelength using the provisions in the header of the SONET/SDH frame into which the ATM signal is mapped, as explained above. The better-received signal is then passed on to the ATM matrix. In the other direction, each signal originated by matrix 92 is duplicated by selector 90 to both transmitters in transceivers 76 associated with each channel or wavelength.

In an IP application, the functionality of the protocol and tributary subsystems 58, 60 resembles that of an IP router. In this case, matrix 92 is an IP routing matrix that has N input port and N corresponding output ports. An IP channel (SONET/SDH framed) is present on each input port at one of the rates supported by matrix 92. A channel of the same rate is output by matrix 92 on the corresponding output port. Matrix 92 routes IP packets from incoming channels to output channels.

The routing for a packet is determined by the contents in its header (i.e., the destination IP address).

Again, the input/output ports are associated with either tributary cards 88 or WDM transceivers 76 that handle channels originating from terminal nodes 36 on virtual star network 20. Matrix 92 allows a packet originating in any terminal nodes 26 or local client systems coupled to tributary cards 88 to be transmitted to any other terminal node 26 or local client system. Matrix 92 contains buffering and routing resources as is known in the art. Electronic components which allow the design of an IP routing matrix are commercially available from various vendors such as MMC Networks and can be used for matrix 92 in an IP application.

Selector 90 in IP applications monitors the quality of the two received signals per wavelength using the provisions in the header of the SONET/SDH frame into which IP packet is mapped similar to applications discussed above. The better-received signal is then passed on to the IP routing matrix. Likewise, in the other direction, each signal originated by matrix 92 is duplicated by selector 90 to both transmitters in transceivers 76 associated with each channel or wavelength.

An alternative head-end node 124 supports a mixed protocol environment as shown in Fig. 12. In mixed protocol environments the requirement is to support multiple traffic types within the same network. An exemplary mixed protocol network includes some SONET/SDH terminals, some ATM terminals, and some IP terminals (all as defined in the previous sections). All signals, including the ATM and IP signals, use SONET/SDH framing.

The basic approach in head-end node 124 is that, given a common SONET/SDH framing used by all signals, a SONET/SDH cross-connect matrix 192 is used to funnel all ATM channels into an ATM switching matrix 194 and all IP traffic into an IP routing matrix 193. SONET/SDH cross-connect matrix 192, ATM switching matrix 194, and IP routing matrix 193 all have the same functionality and structure as described above for matrix 92 in SONET/SDH, ATM, and IP applications, respectively. Selector 190 provides the same functionality as selector 90 discussed above.

Head-end node 192 also includes the same optical subsystem 56 and tributary subsystem 60 design as discussed above. All traffic from tributary cards 88

and from terminal nodes on virtual star network 20 is routed to SONET/SDH matrix 192, which concentrates all SONET/SDH channels carrying ATM traffic to output ports connected to ATM switching matrix 194 and all channels carrying IP traffic to output ports connected to IP routing matrix 193. ATM matrix 194 performs all ATM switching required between these ATM channels as described above and then returns the traffic to the SONET/SDH matrix, which cross-connects them back to the destination tributary ports or terminals. IP traffic channels are handled similarly through IP routing matrix 193. "Regular" SONET/SDH channels requiring only direct SONET/SDH cross-connecting are handled directly by SONET/SDH cross-connect matrix 192.

Another alternative embodiment head-end node 224 as shown in Fig. 13 supports an alternative mixed protocol environment, in which the main protocol used is ATM, and in which some of the ATM virtual connections carry IP traffic. In this case, an ATM switching matrix 294 directly handles any non-IP ATM traffic and channels IP traffic embedded within the ATM traffic to an IP routing matrix 293. Routed IP traffic is then coupled back through ATM switching matrix 294.

It is understood that other routing matrix configurations can be used with or added to the illustrated embodiments. For example, a SONET/SDH cross-connect matrix (not shown) coupled to an IP Routing matrix to accommodate IP traffic not embedded in ATM can be added to the embodiment of Fig. 13. Similarly, other protocol modules can be provided as desired for other communication protocols.

Although the invention has been described in detail with reference to certain preferred embodiments, variations and modifications exist within the scope and spirit of the present invention as described and defined in the following claims.